NEXALUS
COOLING THE CLOUD

Why Liquid Cooling?

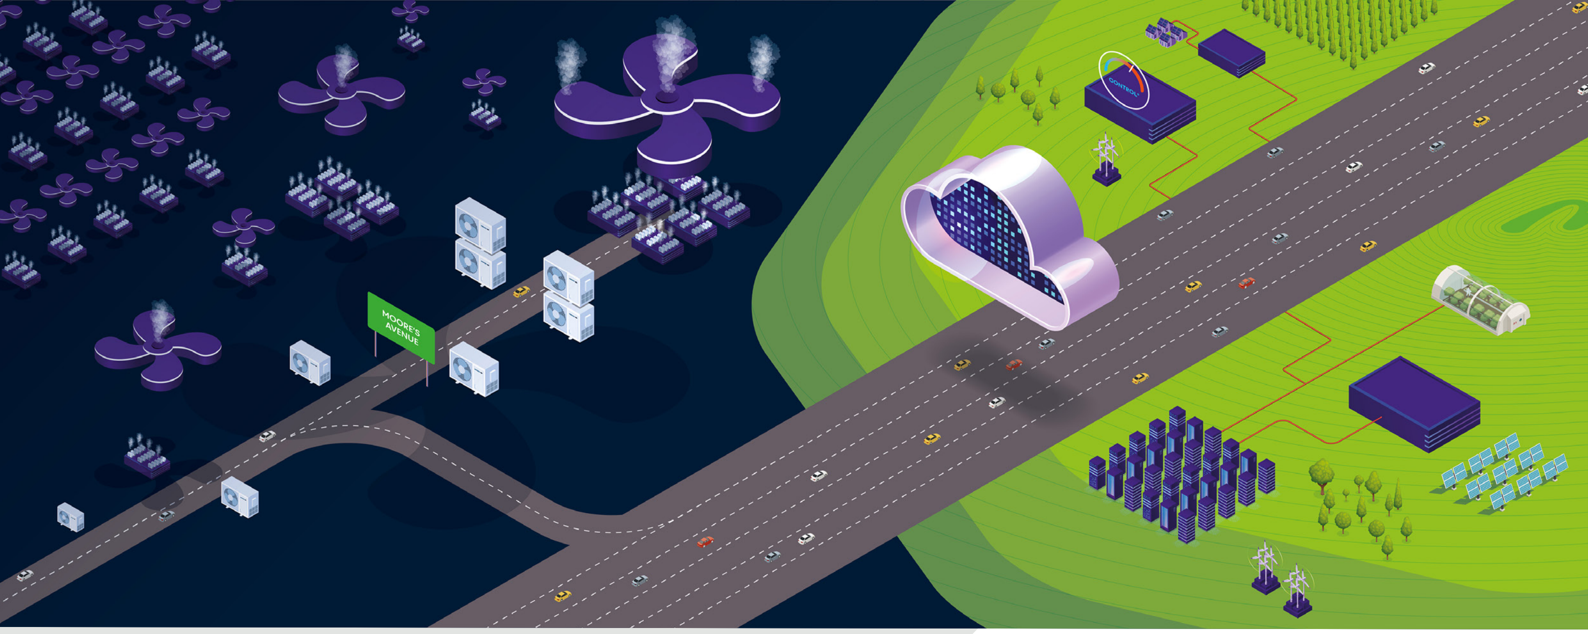

Why Liquid Cooling?

We have reached the technology pivot from air to industrial scale board-level liquid cooling of high-heat flux electronics. This has been on the horizon for a decade or more, and we have watched it getting closer as die-level heat fluxes and package densities relentlessly rose.

THE SIMPLE FACTS ARE:

- Forcing more air does not help since convective heat transfer coefficients have already flat-lined and fan speed, pressure, power and noise levels have hit thresholds.
- The only option available to keep hanging on to board-level air-cooling is to engineer increasingly larger fin structures with bigger/more fans.
- Larger air-cooled heat sinks require sophisticated integrated heat spreaders, such as heat pipes and vapour chambers, as the metal fins get progressively further away from the source.

This ballooning of the heat sink has managed to keep air-cooling on the table. However, the idea of cooling highly integrated, highly miniaturized, highly compact, highly cost-optimized boards with proportionately massive, heavy and expensive heat sinks is contrary with the entire notion of miniaturizing electronics. From an electronics packaging perspective, the board and the air-cooling technologies are moving in opposite directions.

Not only is this technology disparity gap widening, in many cases there is simply no concept of board-level air-cooling that achieves the thermal resistances required to reach limiting junction temperatures. There is a growing subset of CPUs, GPUs, IGBTs, LEDs, PAs etc. that are at a point of no return from liquid cooling. For many others, the go-to technological fix for facilitating air-cooling is to under-power the electronic devices. This is as true for computer CPUs as it is for automotive LED headlamps. This throttling back to an air-cooled thermal design power is inconsistent to the entire point of pushing the envelope of electronics miniaturization and performance.

The ingenuity and innovativeness that has created engineering solutions that were able to keep board-level air-cooling at pace with the heat flux demands of these high-powered electronic components and packages has been impressive. However, we have known for some time that there was a natural end to cooling them with air and that the shift to liquid cooling was inevitable. Like it or not, this time is now.

What is 'high performance' liquid cooling?

One part of 'high performance' relates to the effectiveness of cooling associated with the liquid-cooled heat exchanger. Effective heat transfer coefficients that are an order of magnitude (10X) better than the highest performing board-level air-cooled solution are high-performance. The best in class heat pipe based air-cooled fan-fin heat exchanger offers circa 5,000 W/m²K[1] of cooling effectiveness, meaning that high performance liquid-cooled heat exchangers should achieve in the region of 50,000 W/m²K to be high-performance.

In achieving this level of cooling effectiveness, single-phase liquid water must be used as the working fluid. Though this might be achievable with convective boiling of some dielectric fluids (refrigerants), it would require highly engineered boiling heat exchangers and flow loops, and would be at the outer edge of their heat transfer capability. Furthermore, these fluids suffer from catastrophic boiling crisis events at relatively low heat fluxes, not to mention severe flow instability problems. Two phase cooling technology is simply not ready nor is it capable of keeping up with escalating heat fluxes.

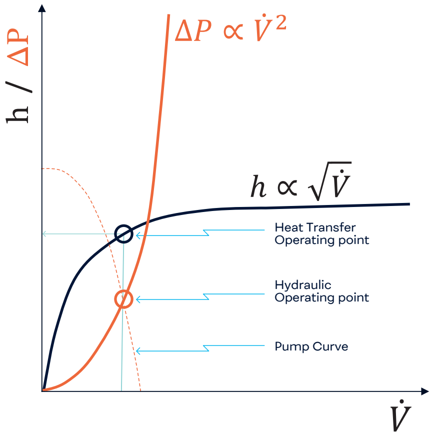

Considering single-phase water-cooling, one must realize that packaging and integration demands that the architecture of the heat sink be commensurate with the electronic components/packages being cooled. Thus, the heat exchanger must scale with the size of the source. Fortuitously, water-cooling effectiveness, as quantified by the convective heat transfer coefficient (*h*), scales with the inverse of the channel dimension, $h \propto 1/d$, which is a positive tension between scaling down the size of the heat exchanger and driving up the heat transfer effectiveness. This positive tension facilitates the engineering of high performance (~50,000 W/m²K) water-cooled heat exchangers with low board-level profiles. It is worth noting here that it is unquestioned scientific knowledge that the highest achievable single-phase convective heat transfer coefficients are created by impinging water microjets.

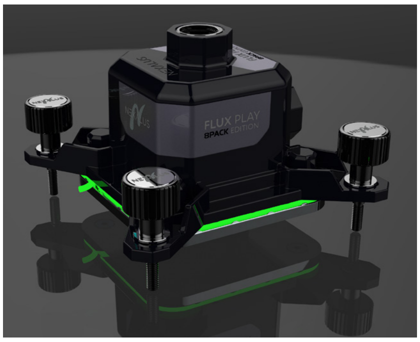

The other part of 'high performance' relates to the hydraulics. A delicate balance must be achieved between pressure drop (ΔP), flow rate ($\dot{V}$) and the resulting heat transfer effectiveness (*h*). This is particularly challenging in all-in-one (AOI) closed pumped systems, where the entire system must be tuned very intentionally to provide a flow rate that creates high cooling levels without extreme pumping pressure. For high-end gaming PC CPU cooling, circa $\dot{V}$~5 L/min with $\Delta P \leq 50$ kPa is acceptable, provided *h*~50,000 W/m²K; this is high performance. Data centres, on the other hand, cannot afford such high flow rates due to the overwhelming demand it would put on the infrastructure in supplying so many servers. Here, the push-pull relationship between the hydraulics and the heat transfer is more severe, requiring much lower flow rates (~5X to 10X lower) and, compared with high-end gaming, a different level of tuning in terms of the hydraulics, heat transfer and thermal design power.

Clearly, there is no single definition of 'high-performance' and there is no silver-bullet solution to liquid cooling across all the emerging fields. The right solution is one that is platform, yet easily tuneable to the thermal-hydraulic design challenge being faced.

1. This is around 3X higher than a conventional direct attach fan-fin metal heat exchanger

How do we keep keeping up?

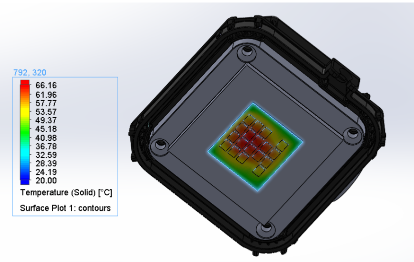

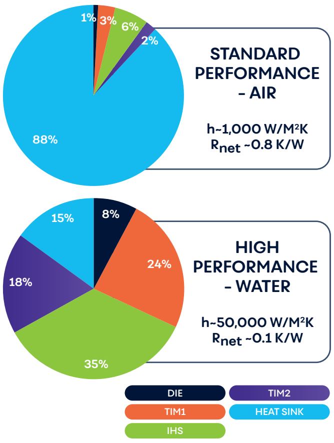

Intel i9 7980E with good TIM1 (0.1 cm²K/W) & "so-so' TIM2 (0.25 cm²K/W) thermal grease

Something interesting happens when you implement high performance liquid cooling in a thermal stack. The thermal stack, given here for an Intel i9 7980XE for an illustrative example, includes the heat source (cores), the semiconductor die, TIM1, the lid (IHS), TIM2 and the attached heat exchanger. When a reasonably good air-cooled fan-fin heat exchanger is used, the overwhelming proportion of the thermal budget is associated with the air-side cooling. From a thermal design standpoint, there is little to be gained by engineering improvements on anything but the air-side. However, when high performance liquid cooling is implemented, not only does the net thermal resistance drop considerably, the relative proportion of the thermal budget associated with each level in the stack is entirely redistributed.

This has the effect of creating a new design paradigm for further advancement in electronics cooling; the next generation of cooling technology will only be realized when the thermal design is approached holistically from source-to-sink. This fully integrated thermal design must consider the entire stack as a single problem and simultaneously address the difficult and physically coupled problems of forced convection, contact resistance, heat spreading, and conjugate heat transfer. This is quite apparent from the example here, where a 50X increase in the attached heat sink effectiveness results in less than 10X improvement in the overall source to sink thermal resistance, with nearly 70% of the thermal budget now being associated with the CPU package itself.

High performance liquid cooling is a game-changer in the context of bringing about a step change in source-to-sink thermal resistance. This effectively unleashes the capability of the electronics while at the same time opens new opportunities in the context of thermal energy recovery and reuse. However, it creates a new type of problem where pushing the envelope even further requires a deeply integrated design process.

Why Nexalus, and why are we different?

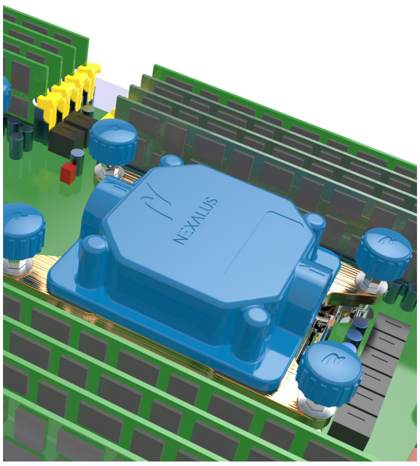

If one scrutinizes today's board-level water cooling technology, it would reveal that once the bells-and-whistles are stripped away, they are really just miniature versions of the air-cooled concept; the metal fins are just smaller, the fluid forced over them is just water, and they are one-size fits all blanket-cooling concepts. In fact, apart from stylizing, there is very little that technically differentiates the leading waterblock brands available in today's market. There is one somewhat hidden reason why they are all ostensibly the same; the microchannel concept they all deploy has converged to its thermal-hydraulic optimum. There is very little, if any, road left for this cooling concept, and this is a mathematical certainty.

At Nexalus, we do not see any sense in merging the lanes of liquid cooling and high-heat flux electronics when one road is near its end and the other has no end in sight. Through diligent research and scientific experience and expertise, we have engineered a new liquid cooling platform that is high performance, yet at the same time (i) is versatile in the context of the broad playing field of today's components and packages, and (ii) is agile in the context of leaving significant thermal-hydraulic headroom for new technology pipelines.

The Nexalus water cooling technology implements jet array impingement heat transfer. This is motivated by the fact that jet impingement convective heat transfer has no performance equal. However, this is only one part of the story. Conventional microchannel cooling must leverage

Why Nexalus, and why are we different?

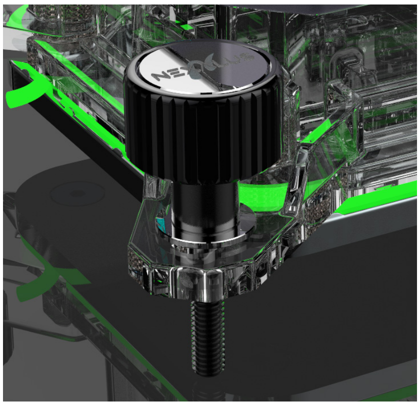

the heightened convective heat transfer coefficients in the small channels fabricated on the surface, whilst simultaneously leveraging the extended surface area the channel walls create; they unanimously focus on the thermal-hydraulics at the heat transfer surface. Conversely, jet array impingement focusses on the fluid mechanics associated with the delivery of the liquid to the surface. By creating very high velocity jets in strategic arrays, extreme levels of cooling are generated when they impinge on the heated surface. This is a completely different way of thinking and, by flipping the engineering from surface structuring to flow delivery, a new thermal-hydraulic design parameter space opens up. This creates an untapped ecosystem for thermal-hydraulic engineering design, including, though not limited to;

Targeted Cooling: From a hydraulic standpoint, the volume of flow, pressure drop and pumping power are not unlimited. Quite the opposite; in practice the available hydraulic resources are quite constrained and are not the same across all platforms. In this context, blanket cooling is a poor design approach since it invariably wastes coolant potential in regions where it is not critical. Jet array impingement heat transfer technology easily allows for focussing the coolant where it is needed. This is as important for concentrated heat sources on single dies as it is for multi-chip modules.

Bespoke Thermal-Hydraulic Design: Cooling of a stack is not a convection problem; it is a conjugate problem of convection and diffusion. Localized heat sources conduct and spread heat through the die which itself conducts and spreads heat through the lid/IHS. This heat diffusion problem is strongly coupled with the convection-side heat transfer problem (termed conjugate heat transfer). By simple variations in the jet size and spacing, distributed convective cooling can be engineered in such a way that the interplay between convection and diffusion is leveraged to double-down on the overall heat transfer effectiveness of the stack, providing bespoke source-to-sink thermal-hydraulic optimization.

Hybrid Cooling: Jet array impingement water-cooling is already superior to its microchannel counterpart. This is before any enhancement features are engineered on the target surface to enhance the jet impingement heat transfer. Jet cooling is unique in this aspect since hybrid cooling, being that where the jet impingement heat transfer is augmented by surface structuration, is a card that has not been necessary to play yet. Furthermore, a unique hydraulic attribute of jets is that almost all of the required pressure drop occurs across the orifice plate, in creating the high velocity jets. Those associated with the inlet, spent flow channel and outlet manifold are in fact very minor. This means that the addition of heat transfer enhancing surface structures has a disproportionately high enhancement effect on the heat transfer compared with the minor hydraulic penalty they incur. As such, the technology is to some degree future-proof, in that there is significant potential to increase cooling effectiveness as heat fluxes continue to escalate, and this can be achieved without compromising the hydraulic constraints.

The answer to the question: Why Nexalus and why are we different? is this - Nexalus is a thermal-fluids science company and we understand the complex physics before we engineer the technological solution. This has resulted in the highest performing and most agile family of water-cooled CPU and graphics card cooling systems on the market, and we are just getting started.

www.nexalus.com

CONTACT US:

Nexalus Ltd,
Unit 13, South Bank,
Crosse's Green,
Cork, Ireland,
T12 XT71

Nexalus Labs,
Water St,
Loughanalla,
Castlepollard,
Co. Westmeath, Ireland,
N91 EP90

Nexalus Research,
Trinity Research & Innovation,
O'Reilly Institute, Trinity College,
Dublin 2, Ireland,

info@nexalus.com
www.nexalus.com

Registered in Ireland No. 628880,
VAT No. IE3573607SH

WITH ENDORSEMENTS AND SUPPORT FROM:

ENTERPRISE IRELAND where innovation means business
Science Foundation Ireland sfi For what's next
Trinity College Dublin Coláiste na Tríonóide, Baile Átha Cliath The University of Dublin
CONNECT
IRISH RESEARCH COUNCIL An Chomhairle um Thaighde in Éirinn

QUALITY CERTIFICATIONS:

IATF 16949 / Automotive Standard
ISO 13485 / Medical Device Standard
ISO 50001 / Energy Management Standard
ISO 45001 / Occupational Heath & Safety Standard
ISO 14001 / Environmental Standard

NEXALUS COOLING THE CLOUD